| Ref # | Hits | Search Query | DBs | Default Operator | Plurals | Time Stamp |
|----------|------|---|----------------------------|---------------------|---------|------------------|
| L2 | 333 | (715/511).CCLS. | US-PGPUB; USPAT | OR | OFF | 2005/09/07 14:38 |
| L4 | 946 | (715/530).CCLS. | US-PGPUB; USPAT | OR | OFF | 2005/09/07 15:27 |
| L5 | 153 | (715/509).CCLS. | US-PGPUB; USPAT | OR | OFF | 2005/09/07 14:39 |
| L6 | 2474 | associat\$4 with version with data | US-PGPUB; USPAT | OR | ON | 2005/09/07 15:27 |
| L7 | 24 | stor\$4 with recent with version with table | US-PGPUB; USPAT | OR | ON | 2005/09/07 15:28 |
| L8 | 3 | 6 and 7 | US-PGPUB; USPAT | OR | ON | 2005/09/07 15:28 |
| L9 | 1309 | stor\$4 with version with table | US-PGPUB; USPAT | OR | ON | 2005/09/07 15:28 |
| L10 | 161 | 6 and 9 | US-PGPUB; USPAT | OR | ON | 2005/09/07 15:28 |
| L12 | 99 | stor\$4 with version with second with table | US-PGPUB; USPAT | OR | ON | 2005/09/07 15:28 |
| L13 | 20 | 10 and 12 | US-PGPUB; USPAT | OR | ON | 2005/09/07 15:28 |
| L14 | 1 | 8 and 13 | US-PGPUB; USPAT | OR | ON | 2005/09/07 15:28 |
| L15 | 58 | determin\$4 with version with number with location | US-PGPUB; USPAT | OR | ON | 2005/09/07 15:29 |
| L16 | 1 | 13 and 15 | US-PGPUB; USPAT | OR | ON | 2005/09/07 15:29 |
| L17 | 57 | 15 not 16 | US-PGPUB; USPAT | OR | ON | 2005/09/07 15:29 |
| L18 | 197 | version with (different near value) | US-PGPUB; USPAT | OR | ON | 2005/09/07 15:30 |
| L19 | 11 | 18 and 9 | US-PGPUB; USPAT | OR | ON | 2005/09/07 15:30 |
| L20 | 10 | 19 not 16 | US-PGPUB; USPAT | OR | ON | 2005/09/07 15:30 |
| L21 | 10 | increment\$4 with maximum with value with version | US-PGPUB; USPAT | OR | ON | 2005/09/07 15:30 |
| L22 | 1 | sequential adj query adj language adj description adj area | US-PGPUB; USPAT | OR | ON | 2005/09/07 15:31 |
| L24 | 456 | old\$4 with version with (table or database) | US-PGPUB; USPAT; EPO | OR | ON | 2005/09/07 15:31 |
| L25 | 1749 | new\$4 with version with (table or database) | US-PGPUB; USPAT; EPO | OR | ON | 2005/09/07 15:31 |

| L26 | 290 | 24 and 25 | US-PGPUB; USPAT; EPO | OR | ON | 2005/09/07 15:31 |
|-----|-----|-----------|----------------------------|----|----|------------------|
| L27 | 1 | 18 and 26 | US-PGPUB; USPAT; EPO | OR | ON | 2005/09/07 15:31 |

| | Document ID | Issue Date | Title | Current OR |
|---|-------------------------|------------|---|------------|
| 1 | US 20050193015 A1 | 20050901 | Method and apparatus for organizing, sorting and navigating multimedia content | 707/104.1 |
| 2 | US 20050081107 A1 | 20050414 | Method and system for autonomic execution path selection in an application | 714/38 |
| 3 | US 20050081019 A1 | | Method and system for autonomic monitoring of semaphore operation in an application | 712/227 |
| 4 | US 20050081010 A1 | 20050414 | Method and system for autonomic performance improvements in an application via memory relocation | 711/165 |
| 5 | US 20050065879 A1 | 20050324 | System and method for web service billing | 705/40 |
| 6 | US 20050049916 A1 | 20050303 | Method and system for facilitating client driven customer and employee performance improvement promotions | 705/14 |
| 7 | US 20040259062 A1 | 20041223 | Method and apparatus for enhancing the integrity of mental competency tests | 434/236 |

| | Document ID | Issue Date | Title | Current OR |
|----|-------------------------|------------|--|------------|
| 8 | US 20040157254 A1 | 20040812 | System and method for designing probes using heterogeneous genetic information, and computer readable medium | 435/6 |
| 9 | US 20040156371 A1 | 20040812 | High speed parser | 370/401 |
| 10 | US 20040155695 A1 | 20040812 | Multiplexer | 327/407 |
| 11 | US 20040081996 A1 | 20040429 | Methods and products related to genotyping and DNA analysis | 435/6 |
| 12 | US 20030184475 A1 | 20031002 | Radio-frequency badge for location measurement | 342/465 |
| 13 | US 20030070087 A1 | 20030410 | System and method for automatic updating of multiple anti-virus programs | 713/201 |
| 14 | US 20020181369 A1 | 20021205 | Copy protection system and method for optical disks and a copy protected optical disk | 369/53.21 |
| 15 | US 20020138452 A1 | 20020926 | Electronic device with automatic capability for location-specific software configuration | 705/401 |
| 16 | US 20010053174 A1 | 20011220 | Spread spectrum localizers | 375/130 |
| 17 | US 20010033607 A1 | 20011025 | Spread spectrum localizers | 375/150 |

| | Do | cument ID | Issue Date | Title | Current OR |
|----|----------|-----------|------------|--|------------|
| 18 | US B1 | 6877037 | 20050405 | Method of updating client's installed data in response to a user-triggered event | 709/227 |
| 19 | US B1 | 6850980 | 20050201 | Content routing service protocol | 709/226 |
| 20 | US B2 | 6825706 | 20041130 | Multiplexer | 327/407 |
| 21 | US B1 | 6810519 | 20041026 | Achieving tight binding for dynamically loaded software modules via intermodule copying | 717/166 |
| 22 | US B1 | 6801331 | 20041005 | Method and system for controlling and communicating with machines using multiple communication formats | 358/1.15 |
| 23 | US B2 | 6795491 | 20040921 | Spread spectrum localizers | 375/213 |
| 24 | US B2 | 6720922 | 20040413 | Radio-frequency badge for location measurement | 342/465 |
| 25 | US B1 | 6708333 | 20040316 | Method and system for reporting failures of a program module in a corporate environment | 717/171 |
| 26 | US B1 | 6703228 | 20040309 | Methods and products related to genotyping and DNA analysis | 435/91.2 |
| 27 | US B1 | 6665797 | 20031216 | Protection of software again against unauthorized use | 713/193 |

| | Do | cument ID | Issue Date | Title | Current OR |
|----|----------|-----------|------------|---|------------|
| 28 | US B2 | 6621781 | 20030916 | Copy protection system and method for optical disks and a copy protected optical disk | 369/47.53 |
| 29 | US B1 | 6545978 | 20030408 | Network managing method and system | 370/236.2 |
| 30 | US B1 | 6544179 | 20030408 | Ultrasound imaging system and method having automatically selected transmit focal positions | 600/447 |
| 31 | US B1 | 6456933 | 20020924 | Navigation method and navigation system for means of locomotion | 701/209 |
| 32 | US B1 | 6452885 | 20020917 | Copy protection system and method for optical disks and a copy protected optical disk | 369/53.21 |
| 33 | US B2 | 6400754 | 20020604 | Spread spectrum localizers | 375/140 |
| 34 | US B1 | 6385268 | 20020507 | Spread spectrum localizers | 375/377 |
| 35 | US B1 | 6330628 | 20011211 | Memory including a data structure used for identifying a communication protocol of a received communication | 710/105 |
| 36 | US B1 | 6330566 | 20011211 | Apparatus and method for optimizing client- state data storage | 707/104.1 |
| 37 | US B1 | 6282112 | 20010828 | Network storage system | 365/52 |

| | Do | cument I | D Issue Date | Title | Current OR |
|----|----------|----------|--------------|---|------------|
| 38 | US B1 | 6275981 | 20010814 | Method and system for correlating profile data dynamically generated from an optimized executable program with source code statements | 717/158 |
| 39 | US B1 | 6240369 | 20010529 | Transmitting location-specific weather-related data to terminals within a plurality of regions | 702/3 |
| 40 | US B1 | 6188973 | 20010213 | Automatic mapping, monitoring, and control of computer room components | 702/188 |
| 41 | US B1 | 6182286 | 20010130 | Dynamic versioning system for multiple users of multi-module software systems | 717/122 |
| 42 | US B1 | 6182284 | 20010130 | Method and system for eliminating phi instruction resource interferences and redundant copy instructions from static-single-assignment-form computer code | 717/146 |
| 43 | US B1 | 6178470 | 20010123 | Chip for CCSDS compatible serial data streams | 710/52 |
| 44 | US | 6041319 | A 20000321 | Method and system for telephone updates of postal scales | 705/409 |
| 45 | US | 6002708 | A 19991214 | Spread spectrum localizers | 375/130 |

| | Do | cument | ID | Issue I | Date | Title | Current | OR |
|----|----|---------------------|-----|---------|------|--|----------|----|
| 46 | US | 5956665 | 5 A | 1999092 | 1 | Automatic mapping, monitoring, and control of computer room components | 702/188 | |
| 47 | US | 595479 ⁻ | 7 A | 1999092 | 1 | System and method for maintaining compatibility among network nodes connected to a computer network | 709/223 | |
| 48 | US | 5881292 | 2 A | 1999030 | 9 | Dynamic versioning system for multiple users of multi-module software system | 717/170 | |
| 49 | US | 5875200 |) A | 1999022 | 3 | Reed-Solomon code system employing k- bit serial techniques for encoding and burst error trapping | 714/784 | |
| 50 | US | 5818603 | 3 A | 1998100 | 6 | Method and system for controlling and communicating with machines using multiple communication formats | 358/296 | |
| 51 | US | 5757372 | 2 A | 1998052 | 'n | Multiple nonlinear undo buttons | 715/840 | |
| 52 | US | 5748891 | A | 1998050 | 5 | Spread spectrum localizers | 375/150 | |
| 53 | US | 5467134 | l A | 1995111 | 4 | Method and system for compressing video data | 348/409. | 1 |
| 54 | US | 4832486 | 5 A | 1989052 | 3 | Method and apparatus or in vitro evalution of focal length and focal length changes in lenses from human and animal eyes | 356/125 | |

| | Document ID | Issue Date | Title | Current OR |
|----|--------------|------------|--|------------|
| 55 | US 4714996 A | 19871222 | Impact calculation for version management in a distributed information service | 707/203 |
| 56 | US 4376977 A | 19830315 | Computer system with scannable program memory | 711/201 |
| 57 | US 3701972 A | 19721031 | DATA PROCESSING SYSTEM | 710/36 |



Subscribe (Full Service) Register (Limited Service, Free) Login

Search: • The ACM Digital Library O The Guide

version data store table

THE ACM DIGITAL LIBRARY

Feedback Report a problem Satisfaction survev

Terms used version data store table

Found 94,390 of 160,906

Sort results Display

results

Best 200 shown

relevance

expanded form

Save results to a Binder Search Tips

Open results in a new

Try an Advanced Search Try this search in The ACM Guide

window

٠

Results 1 - 20 of 200

Result page: **1** <u>2</u> <u>3</u> <u>4</u> <u>5</u> <u>6</u> <u>7</u> <u>8</u> <u>9</u> <u>10</u>

next

Relevance scale 🔲 📟 📟

1 Query and view processing: On querying versions of multiversion data warehouse Tadeusz Morzy, Robert Wrembel

November 2004 Proceedings of the 7th ACM international workshop on Data warehousing and OLAP

Full text available: pdf(244.86 KB) Additional Information: full citation, abstract, references, index terms

A data warehouse (DW) is fed with data that come from external data sources that are production systems. External data sources, which are usually autonomous, often change not only their content but also their structure. The evolution of external data sources has to be reflected in a DW, that uses the sources. Traditional DW systems offer a limited support for handling dynamics in their structure and content. A promising approach to handling changes in DW structure and content is based on a mu ...

Keywords: data versioning, data warehouse, metadata, multiversion query, schema versioning

2 Database theory, technology and applications (DTTA): Creation and management of versions in multiversion data warehouse



Bartosz B□bel, Johann Eder, Christian Koncilia, Tadeusz Morzy, Robert Wrembel March 2004 Proceedings of the 2004 ACM symposium on Applied computing

Full text available: pdf(516.99 KB)

Additional Information: full citation, abstract, references, citings, index

A data warehouse (DW) provides an information for analytical processing, decision making, and data mining tools. On the one hand, the structure and content of a data warehouse reflects a real world, i.e. data stored in a DW come from real production systems. On the other hand, a DW and its tools may be used for predicting trends and simulating a virtual business scenarios. This activity is often called the what-if analysis. Traditional DW systems have static structure of their schemas and relati ...

Keywords: data warehouse, integrity constraints, versioning

3 Industrial sessions: beyond relational tables: Coordinating backup/recovery and data consistency between database and file systems Suparna Bhattacharya, C. Mohan, Karen W. Brannon, Inderpal Narang, Hui-I Hsiao,

Mahadevan Subramanian

June 2002 Proceedings of the 2002 ACM SIGMOD international conference on Management of data

Full text available: pdf(1.44 MB) Additional Information: full citation, abstract, references, index terms

Managing a combined store consisting of database data and file data in a robust and consistent manner is a challenge for database systems and content management systems. In such a hybrid system, images, videos, engineering drawings, etc. are stored as files on a file server while meta-data referencing/indexing such files is created and stored in a relational database to take advantage of efficient search. In this paper we describe solutions for two potentially problematic aspects of such a data ...

Keywords: DB2, content management, database backup, database recovery, datalinks

Using Applications of Data Versioning in Database Application Development Ramkrishna Chatterjee, Gopalan Arun, Sanjay Agarwal, Ben Speckhard, Ramesh Vasudevan May 2004 Proceedings of the 26th International Conference on Software Engineering

Full text available: pdf(166.57 KB) Additional Information: full citation, abstract

Database applications such as enterprise resource planning systems and customer relationship management systems are widely used software systems. Development and testing of database applications is difficult because the program execution depends on the persistent state stored in the database. In this paper we show that how versioning of the persistent data stored in the database can solve some critical problems in the development and testing of database applications can be solved by vers ...

⁵ A linear-time scheme for version reconstruction

Lin Yu, Daniel J. Rosenkrantz

May 1994 ACM Transactions on Programming Languages and Systems (TOPLAS), Volume 16 Issue 3

Additional Information: full citation, abstract, references, citings, index Full text available: pdf(1.47 MB) terms, review

An efficient scheme to store and reconstruct versions of sequential files is presented. The reconstruction scheme involves building a data structure representing a complete version, and then successively modifying this data structure by applying a sequence of specially formatted differential files to it. Each application of a differential file produces a representation of an intermediate version, with the final data structure representing the requested version. The scheme uses a ...

Keywords: data structures, database systems, differential files, document preparation, software systems, textual objects, version control

6 Effective fine-grain synchronization for automatically parallelized programs using optimistic synchronization primitives

Martin C. Rinard

November 1999 ACM Transactions on Computer Systems (TOCS), Volume 17 Issue 4

Additional Information: full citation, abstract, references, citings, index Full text available: pdf(637.69 KB) terms, review

This article presents our experience using optimistic synchronization to implement finegrain atomic operations in the context of a parallelizing compiler for irregular, object-based computations. Our experience shows that the synchronization requirements of these programs differ significantly from those of traditional parallel computations, which use loop nests to access dense matrices using affine access functions. In addition to coarse-grain barrier synchronization, our irregular comput ...



Keywords: atomic operations commutativity analysis, optimistic synchronization, parallel computing, parallelizing compilers, synchronization

7 The BUCKY object-relational benchmark



Michael J. Carey, David J. DeWitt, Jeffrey F. Naughton, Mohammad Asgarian, Paul Brown, Johannes E. Gehrke, Dhaval N. Shah

June 1997 ACM SIGMOD Record, Proceedings of the 1997 ACM SIGMOD international conference on Management of data, Volume 26 Issue 2


Full text available: pdf(1.48 MB)

Additional Information: full citation, abstract, references, citings, index terms

According to various trade journals and corporate marketing machines, we are now on the verge of a revolution—the object-relational database revolution. Since we believe that no one should face a revolution without appropriate armaments, this paper presents BUCKY, a new benchmark for object-relational database systems. BUCKY is a query-oriented benchmark that tests many of the key features offered by object-relational systems, including row types and inheritance, references and path e ...

8 SmartFiles: an OO approach to data file interoperability





Volume 30 Issue 10

Full text available: pdf(1.59 MB)

Additional Information: full citation, abstract, references, citings, index terms

Data files for scientific and engineering codes typically consist of a series of raw data values whose description is buried in the programs that interact with these files. In this situation, making even minor changes in the file structure or sharing files between programs (interoperability) can only be done after careful examination of the data files and the I/O statements of the programs interacting with this file. In short, scientific data files lack selfdescription, and other self-describin ...

9 P1: "Yes, but does it scale?": practical considerations for database-driven information systems



John Russell

October 2001 Proceedings of the 19th annual international conference on Computer documentation

Full text available: pdf(231.31 KB)

Additional Information: full citation, abstract, references, citings, index terms

This paper explores the process of designing and implementing a database-driven system of online documentation, and putting it live on the web for customers to use. Using real-life examples, it discusses practical considerations for balancing performance, scalability, and reliability.

Keywords: Oracle, automation, categorization, database, performance, reliability, scalability, web services

10 Change management: An infrastructure for development of object-oriented, multi-level configuration management services

Tien N. Nguyen, Ethan V. Munson, John T. Boyland, Cheng Thao

May 2005 Proceedings of the 27th international conference on Software engineering

Full text available:

Additional Information:

pdf(418.74 KB)

full citation, abstract, references, index terms

In an integrated development environment, the ability to manage the evolution of a software system in terms of logical abstractions, compositions, and their interrelations is crucial to successful software development. This paper presents a novel framework and infrastructure, Molhado, upon which to build object-oriented software configuration management (SCM) services in a SCM-centered integrated development environment. Key contributions of this paper include a product versioni ...

Keywords: software configuration management, version control

11 Profiling a parallel language based on fine-grained communication

Bjoern Haake, Klaus E. Schauser, Chris Scheiman

November 1996 Proceedings of the 1996 ACM/IEEE conference on Supercomputing (CDROM) - Volume 00

Full text available: pdf(198.94 KB) Additional Information: full citation, abstract, references, index terms

Fine tuning the performance of large parallel programs is a very difficult task. A profiling tool can provide detailed insight into the utilization and communication of the different processors, which helps identify performance bottlenecks. In this paper we present a profiler for the fine-grained parallel programming language Split-C, which provides a simple global address space memory model. As our experience shows, it is much more challenging to profile programs that make use of efficient ...

Keywords: Parallel programming, performance analysis, profiling, fine-grained communication, Split-C, Active Messages

12 An experimental analysis of the performance of fourth generation tools on PCs

Victor M. Matos, Paul J. Jalics

November 1989 Communications of the ACM, Volume 32 Issue 11

Full text available: pdf(1.52 MB)

Additional Information: full citation, abstract, references, citings, index terms, review

The performance of several Fourth Generation Language (4GL) tools is analyzed empirically and compared with equivalent programs written in the third generation COBOL programming language. A set of performance benchmarks consisting of thirteen separate functions is presented which encompasses the areas of simulating the operators of the relational algebra, accessing records in the database, and updating the database. This serves as a baseline for comparing the various 4GL systems.

13 Satellite-based information services: Cost based data dissemination in satellite networks

Bo Xu, Ouri Wolfson, Sam Chamberlain, Naphtali Rishe January 2002 Mobile Networks and Applications, Volume 7 Issue 1

Full text available: pdf(312.11 KB)

Additional Information: full citation, abstract, references, citings, index terms

We consider the problem of data dissemination in a broadcast network. In contrast to previously studied models, broadcasting is among peers, rather than client server. Such a model represents, for example, satellite communication among widely distributed nodes, sensor networks, and mobile ad hoc networks. We introduce a cost model for data dissemination in peer to peer broadcast networks. The model quantifies the tradeoff between the inconsistency of the data, and its transmission cost; the tran ...

Keywords: data replication, distributed databases, satellite networks





14 Are bitvectors optimal?

H. Buhrman, P. B. Miltersen, J. Radhakrishnan, S. Venkatesh

May 2000 Proceedings of the thirty-second annual ACM symposium on Theory of computing

Full text available: pdf(1.01 MB)

Additional Information: full citation, references, citings, index terms

15 Database performance in the real world: TPC-D and SAP R/3

Joachen Doppelhammer, Thomas Höppler, Alfons Kemper, Donald Kossmann June 1997 ACM SIGMOD Record, Proceedings of the 1997 ACM SIGMOD international conference on Management of data, Volume 26 Issue 2

Full text available: pdf(1.54 MB)

Additional Information: full citation, abstract, references, citings, index terms

Traditionally, database systems have been evaluated in isolation on the basis of standardized benchmarks (e.g., Wisconsin, TPC-C, TPC-D). We argue that very often such a performance analysis does not reflect the actual use of the DBMSs in the "real world." End users typically don't access a stand-alone database system; rather they use a comprehensive application system, in which the database system constitutes an integrated component. In order to derive performance evalu ...

16 Experiments with an ocean circulation model on CEDAR

L. DeRose, K. Gallivan, E. Gallopoulos

August 1992 Proceedings of the 6th international conference on Supercomputing

Full text available: pdf(1.41 MB)

Additional Information: full citation, abstract, references, index terms

We present the design of the GFDL ocean circulation model as adapted for simulations of the Mediterranean basis for the Cedar multicluster architecture. The model simulates the basic aspects of large-scale, baroclinic ocean circulation, including treatment of irregular bottom topography. The data and computational mapping strategies and their effect on the design are discussed. The code was parametrized to offer several choices for data partitionings of the computational domain, for placeme ...

17 The logical disk: a new approach to improving file systems

Wiebren de Jonge, M. Frans Kaashoek, Wilson C. Hsieh

December 1993 ACM SIGOPS Operating Systems Review, Proceedings of the fourteenth ACM symposium on Operating systems principles, Volume 27 Issue 5

Full text available: pdf(1.55 MB)

Additional Information: full citation, abstract, references, citings, index terms

The Logical Disk (LD) defines a new interface to disk storage that separates file management and disk management by using logical block numbers and block lists. The LD interface is designed to support multiple file systems and to allow multiple implementations, both of which are important given the increasing use of kernels that support multiple operating system personalities. A log-structured implementation of LD (LLD) demonstrates that LD can be implemented efficiently. LLD adds about 5% to 10% ...

Keywords: MINIX, UNIX, disk storage management, file system organization, file system performance, high write performance, log-structured file system, logical disk

18 Physical design: Handling big dimensions in distributed data warehouses using the DWS technique



Marco Costa, Henrique Madeira

November 2004 Proceedings of the 7th ACM international workshop on Data warehousing and OLAP

Full text available: pdf(288.33 KB) Additional Information: full citation, abstract, references, index terms

The DWS (Data Warehouse Striping) technique allows the distribution of large data warehouses through a cluster of computers. The data partitioning approach partition the facts tables through all nodes and replicates the dimension tables. The replication of the dimension tables creates a limitation to the applicability of the DWS technique to data warehouses with big dimensions. This paper proposes a strategy to handle large dimensions in a distributed DWS system and evaluates the proposed str ...

Keywords: data warehousing, distributed query execution

19 Compiler-directed run-time monitoring of program data access

Chen Ding, Yutao Zhong

June 2002 ACM SIGPLAN Notices, Proceedings of the 2002 workshop on Memory system performance, Volume 38 Issue 2 supplement

Full text available: pdf(1.40 MB)

Additional Information: full citation, abstract, references, citings

Accurate run-time analysis has been expensive for complex programs, in part because most methods perform on all a data. Some applications require only partial reorganization. An example of this is off-loading infrequently used data from a mobile device. Complete monitoring is not necessary because not all accesses can reach the displaced data. To support partial monitoring, this paper presents a framework that includes a source-to-source C compiler and a run-time monitor. The compiler inserts ru ...

²⁰ Making data structures confluently persistent

Amos Fiat, Haim Kaplan

January 2001 Proceedings of the twelfth annual ACM-SIAM symposium on Discrete algorithms

Full text available: pdf(924.65 KB)

Additional Information: full citation, abstract, references, citings, index terms

We address a longstanding open problem of [8, 7], and present a general transformation that takes any data structure and transforms it to a confluently persistent data structure. We model this general problem using the concepts of a version DAG (Directed Acyclic Graph) and an instantiation of a version DAG. We introduce the concept of the effective depth of a vertex in the version DAG and use it to derive information theoretic lower bounds on the space expansion of any such transformation for ...

Results 1 - 20 of 200 Result page: **1** 2 3 4 5 6 7 8 9 10 next

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2005 ACM, Inc.

<u>Terms of Usage Privacy Policy Code of Ethics Contact Us</u>

Useful downloads: Adobe Acrobat QuickTime Windows Media Player Real Player